

BDD-Based Cryptanalysis of Keystream Generators

Matthias Krause

Theoretische Informatik, Universität Mannheim, 68131 Mannheim, Germany,
krause@informatik.uni-mannheim.de

Abstract. Many of the keystream generators which are used in practice are LFSR-based in the sense that they produce the keystream according to a rule $y = C(L(x))$, where $L(x)$ denotes an internal linear bitstream, produced by a small number of parallel linear feedback shift registers (LFSRs), and C denotes some nonlinear compression function. We present an $n^{O(1)}2^{(1-\alpha)/(1+\alpha)n}$ time bounded attack, the FBDD-attack, against LFSR-based generators, which computes the secret initial state $x \in \{0, 1\}^n$ from cn consecutive keystream bits, where α denotes the rate of information, which C reveals about the internal bitstream, and c denotes some small constant. The algorithm uses Free Binary Decision Diagrams (FBDDs), a data structure for minimizing and manipulating Boolean functions. The FBDD-attack yields better bounds on the effective key length for several keystream generators of practical use, so a $0.656n$ bound for the self-shrinking generator, a $0.6403n$ bound for the A5/1 generator, used in the GSM standard, a $0.6n$ bound for the E_0 encryption standard in the one level mode, and a $0.8823n$ bound for the two-level E_0 generator used in the Bluetooth wireless LAN system.

1 Introduction

A keystream generator is a finite automaton, E , determined by a set Q of inner states, a state transition function $\delta_E : Q \rightarrow Q$, and an output function $a_E : Q \rightarrow \{0, 1\}$. The usual case is that $Q = \{0, 1\}^n$ for some integer $n \geq 1$, n is called the keylength of E . Starting from an initial state $x^0 \in Q$, in each time unit i , E outputs a key bit $y_i = a_E(x^i)$ and changes the inner state according to $x^{i+1} = \delta_E(x^i)$. For each initial state $x \in \{0, 1\}^n$ we denote by $y = E(x)$ the keystream $y = y_0, y_1, \dots$ produced by E when starting on x .

Keystream generators are designed for fast online encryption of bitstreams which have to pass an insecure channel. A standard application is to ensure the over-the-air privacy of communicating via mobile cellular telephones. A plaintext bit stream p_0, p_1, p_2, \dots is encrypted into a ciphertext bitstream e_0, e_1, e_2, \dots via the rule $e_i = p_i \oplus y_i$, where $y = E(x)$. The legal receiver knows x and can decrypt the bitstream using the same rule. The only secret information is the initial state x , which is exchanged before starting the transmission using a suitable key-exchange protocol. It is usual to make the pessimistic assumption that an attacker knows not only the encrypted bitstream, but even some short piece of

the plaintext, and, therefore, can easily compute some piece of the keystream. Consequently, one necessary criterion for the security of a keystream generator is that there is no feasible way to compute the secret initial state x from $y = E(x)$. Observe that a trivial exhaustive search attack needs time $n^{O(1)}2^n$.

In this paper we suggest a new type of attack against keystream generators, which we will call FBDD-attack, and show that LFSR-based keystream generators are vulnerable against FBDD-attacks. We will call a generator to be LFSR-based if it consists of two components, a linear bitstream generator L which generates for each initial state $x \in \{0, 1\}^n$ a linear bitstream $L(x)$ by one or more parallel LFSRs, and a compressor C which transforms the internal bitstream into the output keystream $y = C(L(x))$. Due to the ease of implementing LFSRs in hardware, and due to the nice pseudorandomness properties of bitstreams generated by maximal length LFSRs, many keystream generators occurring in practice are LFSR-based.

FBDD is the abbreviation for *free binary decision diagrams*, a data structure for representing and manipulating Boolean function, which were introduced by *Gergov* and *Meinel* in [9] and *Sieling* and *Wegener* in [16]. Due to specific algorithmic properties, FBDDs and in particular Ordered BDDs (OBDDs), a special kind of FBDDs, became a very useful tool in the area of automatic hardware verification (see also the paper of *Bryant* [4] who initiated the study of graph-based data structures for Boolean function manipulation). The important properties of FBDDs are that they can be efficiently minimized, that they allow an efficient enumeration of all satisfying assignments, and that the minimal FBDD-size of a Boolean function is not much larger than the number of satisfying assignments. We show that these properties can also be successfully used for cryptanalysis. The problem of finding a secret key x fulfilling $y = E(x)$ for a given encryption algorithm E and a given ciphertext y can be reduced to finding the minimal FBDD P for the decision if x fulfils $y = E(x)$. If the length of y is close to the unicity distance of E then P is small, and x can be efficiently computed from P . The main result of this paper is that the special structure of LFSR-based keystream generators implies a nontrivial dynamic algorithm for computing this FBDD P .

In particular, the weakness of LFSR-based keystream generators is that the compressor C has to produce the keystream in an online manner. For getting a high bitrate, C should use only a small memory, and should consume only a few new internal bits for producing the next output bit. These requirements imply that the decision if an internal bitstream z generates a prefix of a given keystream y via C can be computed by small FBDDs. This allows to compute dynamically a sequence of FBDDs P_m , $m \geq n$, which test a given initial state $x \in \{0, 1\}^n$ whether $C(L_{\leq m}(x))$ is prefix of y , where $L_{\leq m}(x)$ denotes the first m bits of the internal linear bitstream generated via L on x . On average, the solution becomes unique for $m \geq \lceil \alpha^{-1}n \rceil$, where α denotes the rate of information which C reveals about the internal bitstream. The FBDDs P_m are small at the beginning and again small if m approaches $\lceil \alpha^{-1}n \rceil$, and we will show that all intermediate FBDDs have a size of at most $n^{O(1)}2^{(1-\alpha)/(1+\alpha)n}$. For all m the FBDD P_m has

to read the first $\lceil \gamma m \rceil$ bits of the keystream, where γ denotes the best case compression ratio of C . Thus, our algorithm computes the secret initial state x from the first $\lceil \gamma \alpha^{-1} n \rceil$ bits of $y = E(x)$. Observe that $\gamma = \alpha$ if C consumes always the same number of internal bits for producing one output bit, and $\alpha < \gamma$ if not. It holds $\gamma \alpha^{-1} \leq 2.5$ in all our examples. Note that for gaining a high bit-rate, α and γ should be as large as possible. Our results say that a higher bit-rate has to be paid by a loss of security.

One advantage of the FBDD-attack is that it is a *short-keystream* attack, i.e., the number of keybits needed for computing the secret initial state $x \in \{0, 1\}^n$ is at most cn for some small constant $c \geq 1$. We apply the FBDD-attack to some of the keystream generators which are most intensively discussed in the current literature, the A5/1 generator which is used in the GSM standard, the self-shrinking generator, and the E_0 -encryption standard, which is included in the Bluetooth wireless LAN system. For all these ciphers, the FBDD-attack has a better time behaviour than *all* short-keystream attacks known before. In some cases there have been obtained *long-keystream* attacks which have a better time behaviour. They use a time-memory tradeoff technique suggested by Golić [8] and are based on the assumption that a long piece of keystream of length 2^{dn} , $d < 1$ some constant, is available. We give an overview on previous results and the relations to our results in section 7.

The paper is organized as follows. In section 2 we give some formal definitions concerning LFSR-based keystream generators and present the keystream generators which we want to cryptanalyze. In section 3 FBDDs are introduced. In section 4 we derive FBDD-relevant properties of LFSR-based keystream generators. In section 5 we define the relevant parameters of LFSR-based keystream generators and formulate the main result. Our cryptanalysis algorithm is presented in section 6 and is applied to particular generators in section 7. Due to space restrictions we cannot present all proofs and BDD-constructions in this extended abstract. The complete version of this paper can be received from [11].

At the first glance it may seem contradictory that we consider practical ciphers like the A5/1 generator with variable keylength n . But observe that definitions of LFSR-based keystream generators, even if they originally were designed for fixed keylength, can be generalized to variable keylength in a very natural way, simply by considering the internal LFSRs to have variable length. Considering variable key-length n allows to evaluate the security of ciphers in terms of how many polynomial time operations are necessary for breaking the cipher. This 'rough' way of security evaluation is sufficient in our context, since the aim of this paper is to present only the general algorithmic idea of the FBDD-attack, to give a rather rough estimation of the time behaviour, and to show an inherent weakness of LFSR-based generators. For practical implementations of FBDD attacks against real-life generators much more effort has to be invested for making the particular polynomial time operations as efficient as possible, see the discussion in section 8.

2 LFSR-Based Keystream Generators

Let us call a keystream generator to be LFSR-based if the generation rule $y = E(x)$ can be written as $y = C(L(x))$, where L denotes a linear bitstream generator consisting of one or more LFSRs, and $C : \{0, 1\}^* \rightarrow \{0, 1\}^*$ a nonlinear compression function, which transforms the internal linear bitstream $L(x)$ into the nonlinear (compressed) output keystream $y = C(L(x))$.¹ Formally, an n -LFSR is a device which produces a bitstream

$$L(x) = L_0(x), L_1(x), \dots, L_m(x), \dots$$

on the basis of a public string $c = (c_1, \dots, c_n) \in \{0, 1\}^n$, the generator polynomial, and a secret initial state $x = (x_0, \dots, x_{n-1}) \in \{0, 1\}^n$, according to the relation $L_i(x) = x_i$ for $0 \leq i \leq n - 1$ and

$$L_m(x) = c_1 L_{m-1}(x) \oplus c_2 L_{m-2}(x) \oplus \dots \oplus c_n L_{m-n}(x) \tag{1}$$

for $m \geq n$. Observe that for all $m \geq 1$, $L_m(x)$ is a GF(2)-linear Boolean function in x_0, \dots, x_{n-1} which can be easily determined via iteratively applying (1).

A linear bitstream generator L of keylength n is defined to be an algorithm which, for some $k \geq 1$, generates a linear bitstream $L(x)$

$$L(x) = L_0(x), L_1(x), L_2(x), \dots$$

by k parallel LFSRs L^0, \dots, L^{k-1} of keylengths n_0, \dots, n_{k-1} , where $n = n_0 + \dots + n_{k-1}$. The initial states $x \in \{0, 1\}^n$ for L are formed by the initial states $x^r \in \{0, 1\}^{n_r}$, $r = 0, \dots, k - 1$, of L^0, \dots, L^{k-1} . L produces in each time unit $j \geq 0$ the bit $L_j(x)$ according to the rule $L_j(x) = L_s^r(x^r)$, where $r = j \bmod k$ and $s = j \operatorname{div} k$. Observe that for all $j \geq 0$ $L_j(x)$ is a GF(2)-linear function in x .

The motivation for taking LFSRs as building blocks for keystream generators is that they can be easily implemented using n register cells connected by a feedback channel. Moreover, if the generator polynomial is primitive, they produce bit streams with nice pseudorandomness properties (maximal period, good auto correlation and local statistics). See, e.g., the monograph by *Golomb* [10] or the article by *Rueppel* [15] for more about the theory of shift register sequences. Clearly, LFSR-sequences alone do not provide any cryptographic security. Thus, the aim of the compression function $C : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is to destroy the low linear complexity of the internal linear bit stream while preserving its nice pseudorandomness properties. Many keystream generators occurring in practice are LFSR-based in the above sense. We investigate the following types of generators.

One classical construction (see, e.g., [10] and [15]) is to combine k parallel LFSRs L^0, \dots, L^{k-1} with an appropriate connection function $c : \{0, 1\}^k \rightarrow$

¹ C compresses the internal bit-stream in an online manner, i.e., $C(z')$ is prefix of $C(z)$ if z' is prefix of z , for all $z, z' \in \{0, 1\}^*$. This justifies to write $y = C(L(x))$ despite of the fact that $L(x)$ is assumed to be infinitely long.

$\{0, 1\}$. The keystream $y = y_0, y_1, y_2, \dots$ is defined via the rule

$$y_j = c(L_j^0(x^0), \dots, L_j^{k-1}(x^{k-1})), \quad j \geq 0,$$

where x^r denotes the initial state for L^r , for $r = 0 \dots, k-1$. Let us call generators of this type **Connect- k** constructions.

The Self-Shrinking Generator was introduced by *Meier* and *Staffelbach* in [14]. It consists of only one LFSR L . The compression is defined via the shrinking function $shrink : \{0, 1\}^2 \rightarrow \{0, 1, \varepsilon\}$, defined as $shrink(ab) = b$, if $a = 1$, and $shrink(ab) = \varepsilon$, the empty word, otherwise. The shrinking-function is extended to bit-strings of even length r as

$$shrink(z_0z_1 \dots z_{r-1}) = y_0y_1 \dots y_{r/2-1},$$

where $y_i = shrink(z_{2i}z_{2i+1})$ for $i = 0, \dots, r/2 - 1$. For each initial state x for L , the self-shrinking generator produces the keystream y according to $y = shrink(L(x))$.

The E_0 Generator is the keystream generator used in the Bluetooth wireless LAN system [3]. It is defined as $E_0(x) = C(L(x))$, where the linear bitstream generator L of E_0 consists of 4 LFSRs L^0, \dots, L^3 . The compression is organized by a finite automaton M with external input alphabet $\{0, 1, 2\}$, state space $Q = \{0, 1, \dots, 15\}$ and output alphabet $\{0, 1\}$, which is defined by an output function $a : Q \times \{0, 1, 2\} \rightarrow \{0, 1\}$ and a state transition function $\delta : Q \times \{0, 1, 2\} \rightarrow Q$. The exact specification of M is published in [3] but does not matter for our purpose and is therefore omitted.

The compression $C(z) = y = y_0y_1 \dots y_{m-1}$ of an internal bit-stream

$$z = z_0^0z_0^1z_0^2z_0^3z_1^0z_1^1z_1^2z_1^3 \dots z_{m-1}^0z_{m-1}^1z_{m-1}^2z_{m-1}^3$$

is defined as $y_j = a(q_j, s_j) \oplus t_j$, where $s_j = (z_j^0 + z_j^1 + z_j^2 + z_j^3) \text{ div } 2$ and $t_j = (z_j^0 + z_j^1 + z_j^2 + z_j^3) \text{ mod } 2$, for all $0 \leq j \leq m - 1$. The actual inner state is updated in each cycle according to the rule $q_{j+1} = \delta(q_j, s_j)$, where q_0 denotes the initial state of M . In practice, the E_0 generator is used with key length 128, the four LFSRs have lengths 39, 33, 31, 25.

The E_0 Encryption Standard in the two-level mode (for short, E_0^2 -Generator) of keylength n combines two E_0 devices of internal keylength $N \geq n$ in the following way. For $x \in \{0, 1\}^n$ it holds $y = E_0^2(x) = E_0(z)$, where z denotes the prefix of length N of $E_0(u)$, and where

$$u = (x_0, \dots, x_{n-1}, U_n(x) \dots, U_{N-1}(x)).$$

$U_i, i = n \dots N - 1$, are public GF(2)-linear functions in (x_0, \dots, x_{n-1}) . In practice, the string u results from putting n secret bits together with $N - n$ known dummy bits into the LFSRs and running them a certain number of steps. The Bluetooth system uses $N = 128$, and n can be chosen as 8, 16, 32, or 64. The reason for choosing a larger internal key length N is to achieve a larger effective key length in n .

The A5/1 Generator is used in the GSM standard. The definition was discovered by *Briceno et. al.* [5] via reverse engineering. The A5/1 generator of key-length n consists of 3 LFSRs L^0 , L^1 , and L^2 of key-lengthes n_0 , n_1 , and n_2 . In each time step i , the output key bit y_i is the XOR of the actual output bits of the 3 LFSRs. A clock control decides in each timestep which of the 3 LFSRs are shifted, and which not. The clock control takes for all $k \in \{0, 1, 2\}$ a control value c_k from the N_k -th register cell of L^k , and computes the control value $m = \text{maj}_3(c_0, c_1, c_2)$.² LFSR L^k is shifted if $m = c_k$, for $k = 0, 1, 2$. The control positions N_k are fixed and fulfil $N_k \in \{\lceil \frac{n_k}{2} \rceil - 1, \lceil \frac{n_k}{2} \rceil\}$.

This keystream generation rule can be written down in a $y = C(L(x))$ fashion in the following way. Given an internal bitstream $z = (z_0^0, z_0^1, z_0^2, \dots, z_m^0, z_m^1, z_m^2, \dots)$ the keystream $y = C(z)$ is defined as follows. In each timestep, C holds 3 output positions $i[0], i[1], i[2]$ and 3 control positions $j[0], j[1], j[2]$. C outputs $x_{i[0]}^0 \oplus x_{i[1]}^1 \oplus x_{i[2]}^2$, computes the new control value $m = \text{maj}_3(x_{j[0]}^0, x_{j[1]}^1, x_{j[2]}^2)$, and updates the i - and j -values via $i[k] := i[k] + 1$ and $j[k] := j[k] + 1$, for those $k \in \{0, 1, 2\}$ for which $m = x_{j[k]}^k$. The output positions are initialized by 0. The control positions are initialized by N_0, N_1, N_2 . Note that in the GSM standard the A5/1 generator is used with key length 64, the 3 LFSRs have lengthes 19, 22 and 23

3 Binary Decision Diagrams (BDDs)

For m a natural number let X_m denote the set of m Boolean variables $\{x_0, \dots, x_{m-1}\}$. A BDD P over X_m is an acyclic directed graph with inner nodes of outdegree 2, a distinguished inner node of indegree 0, the source, and two sink nodes of outdegree 0, one 0-sink and one 1-sink. All inner nodes, i.e. nodes of outdegree > 0 , are labelled with queries $x_i?$, $0 \leq i \leq m - 1$, and are left by one edge labelled 0 (corresponding to the answer $x_i = 0$) and one edge labelled 1 (corresponding to the answer $x_i = 1$). Each assignment b to the X_m -variables defines a unique computational path in P , which will be called the b -path in P . The b -path starts at the source, answers always b_i on queries $x_i?$ and, thus, leads to a unique sink. The label of this sink is defined to be the output $P(b) \in \{0, 1\}$ of P on input $b \in \{0, 1\}^m$. We denote by $\text{One}(P) \subseteq \{0, 1\}^m$ the set of inputs accepted by P , $\text{One}(P) = \{b \in \{0, 1\}^m; P(b) = 1\}$. Each BDD P over X_m computes a unique function $f : \{0, 1\}^m \rightarrow \{0, 1\}$, by $f(b) = 1 \iff b \in \text{One}(P)$. The size of P , $|P|$, is defined to be the number of inner nodes of P . Two BDDs are called equivalent if they compute the same function.

We call an BDD P to be a *free* binary decision diagram (FBDD) if along each computational path in P each variable occurs at most once. In [9] and [16] it was observed that FBDDs can be efficiently minimized with respect to all equivalent FBDDs which read the input variables in an equivalent order. The equivalence of orders of reading the input variables is expressed by using the notion of graph orderings.

² maj_3 is defined to output $c \in \{0, 1\}$ iff at least 2 of its 3 arguments have value c .

Definition 1. A graph ordering G of X_m is an FBDD over X_m with only one (unlabelled) sink, for which on each path from the root to the sink all m variables occur.

Graph orderings are not designed for computing Boolean functions. Their aim is to define for each assignments $b = (b_0, \dots, b_{m-1})$ to X_m a unique variable ordering $\pi_G(b) = (x_{i_1(b)}, \dots, x_{i_m(b)})$, namely the ordering in which the variables are requested along the unique b -path in G .

Definition 2. An FBDD is called G -driven, for short, G -FBDD, if the ordering in which the variables are requested along the b -path in P respects $\pi_G(b)$, for all assignments b . I.e., there do not exist assignments b , variables x_i and x_j such that x_i is requested above x_j at $\pi_G(b)$, but below x_j at the b -path in P .

A special, extensively studied variant of FBDDs are Ordered Binary Decision Diagrams (OBDDs). An FBDD P is called OBDD with variable ordering π (for short π -OBDD) if all paths in P respect π .

We need the following nice algorithmic properties of graph-driven FBDDs. Let $f, g : \{0, 1\}^m \rightarrow \{0, 1\}$ be Boolean functions, let G be a graph ordering for X_m , and let P and Q be G -driven FBDDs for f and g , respectively.

Property 3.01 *There is an algorithm MIN which computes from P in time $O(|P|)$ the (uniquely defined) minimal G -driven FBDD $\min(P)$ for f .*

Property 3.02 *It holds that $|\min(P)| \leq m|One(P)|$.*

Property 3.03 *There is an algorithm SYNTH which computes in time $O(|P||Q||G|)$ a G -driven FBDD $P \wedge Q$, $|P \wedge Q| \leq |P||Q||G|$, which computes $f \wedge g$.*

Property 3.04 *There is another algorithm SAT which enumerates all elements in $One(P)$ in time $O(|One(P)||P|)$.*

See, e.g., the book by Wegener [17] for a detailed description and analysis of the OBDD- and FBDD-algorithms. FBDDs together with the procedures MIN, SYNTH and SAT will be the basic data structure used in our cryptanalysis. Instead giving explicit examples of FBDDs we refer to the following relation between algorithms and FBDDs, see, e.g. [12]. Let $F \subseteq \{0, 1\}^*$ decision problem written as a sequence of Boolean functions $(F_m : \{0, 1\}^m \rightarrow \{0, 1\})_{m \in \mathbb{N}}$.

Theorem 1. *Suppose that F can be decided by an algorithm A which reads each input bit at most once and uses for all $m \geq 1$ and all inputs of length m at most $s(m)$ memory cells. Then, for all $m \geq 1$, A provides an efficient construction of an FBDD P_m of size $O(m2^{s(m)})$ for F_m . \square*

Sketch of a Proof. The inner nodes of P_m are organized in m disjoint levels. Each level j , $1 \leq j \leq m$, contains potentially one node v_r^j for each possible assignment $r \in \{0, 1\}^{s(m)}$ to the memory cells used by A on inputs of length m . This implies the size bound claimed in the Theorem. P_m is constructed top down. The first inner node we insert to P_m is the root $v_{r_0}^1$ labelled by x_{i_1} , where r_0 denotes the initial assignment of the memory cells and i_1 the position of the first input bit read by A . If we have inserted an inner node $v = v_r^j$, labelled by x_i , to P_m than for all $b \in \{0, 1\}$ the b -successor v' of v is constructed as follows. If A starting on configuration r after reading $x_i = b$ stops accepting (resp. rejecting) without reading another input bit then v' is defined to be the 1-sink (resp. the 0-sink). Otherwise let r' denote the first configuration when A is going to read a new input bit, say x_k . Then v' is defined to be $v_{r'}^{j+1}$ labelled by x_k . \square

4 FBDD-Aspects of Key-Stream Generators

Let E be a LFSR-based keystream generator of key-length n with linear keystream generator L and compression function $C : \{0, 1\}^* \rightarrow \{0, 1\}^*$. Let $x \in \{0, 1\}^n$ denote an initial state for L .

Definition 3. For all $m \geq 1$ let G_m^C denote the graph ordering, which assigns to each internal bitstream z the order in which C reads the first m bits of z .

Observe that for the E_0 generator, the self-shrinking generator, as well as for Connect- k generators, the order in which the compressor reads the internal bits does not depend of the internal bitstream itself, i.e., G_m^C has size m and G_m^C -driven FBDDs are OBDDs. But in the case of the A5/1 generator, this order is governed by the clock control, and can be different for different inputs. The efficiency of our cryptanalysis algorithm is based on the following FBDD assumption on E .

FBDD Assumption. The graph ordering G_m^C has polynomial size in m . Moreover, for arbitrary keystreams y , the predicate if for given $z \in \{0, 1\}^m$ $C(z)$ is prefix of y can be computed by G_m^C -FBDDs of polynomial size in m .

It is quite easy to see that the compression function of a Connect- k generators, defined by a function $c : \{0, 1\}^k \rightarrow \{0, 1\}$, fulfils the FBDD-assumption. The compressor reads the internal bits in the canonical order $\pi = 0, 1, 2, 3, \dots$. Linear size π -OBDDs which decide whether $z \in \{0, 1\}^m$ generates the first $\lfloor m/k \rfloor$ bits of a given keystream y via c can be constructed, via Theorem 1, according to the following algorithm.

1. For $j := 0$ to $\lfloor m/k \rfloor$
2. if $c(z_j^0, \dots, z_j^{k-1}) \neq y_j$ then stop(0)
3. stop(1)

Quadratic size π -OBDDs which decide for $z \in \{0, 1\}^m$ whether $shrink(z)$ is prefix of a given keystream y can be constructed, via Theorem 1, according to the following algorithm

1. $k := 0, j := 0$
2. while $j < m - 1$
3. if $z_j = 0$
4. then $j := j + 2$
5. else
6. if $z_{j+1} = y_k$
7. then $j := j + 2, k := k + 1$
7. else stop(0)
8. stop(1)

The FBDD constructions for all the E_0 -, the E_0^2 -, and the A5/1 are sketched in section 7, resp. can be studied in the long version [11].

We still need to estimate the size of FBDDs which decide whether a given $z \in \{0, 1\}^m$ is a linear bit-stream.

Lemma 1. *For all $m \geq n$, the decision whether $z \in \{0, 1\}^m$ is generated via linear bitstream generator L of keylength n can be computed by a G_m^C -driven FBDD of size at most $|G_m^C|2^{m-n}$.*

Proof: Let V_m denote the set of inner nodes of G_m^C . We construct a G_m^C -driven FBDD R_m with the set $W_m = V_m \times \{0, 1\}^{m-n}$ of inner nodes.

For all initial states $x \in \{0, 1\}^n$ and all internal positions $j, n \leq j \leq m - 1$, write $L_j(x)$ as

$$L_j(x) = \bigoplus_{k=0}^{n-1} L_{k,j} x_k.$$

G_m^C ensures that x_k is always read before x_j if $L_{k,j} = 1$.

Let the root of R_m be the node $(v_0, \vec{0})$ where v_0 denotes the root of G_m^C . Let all nodes (v, b) have the same label as v does in G_m^C . The edges of R_m are defined according to the following rules. Let $v \in V_m$ and $b = (b_n, \dots, b_{m-1}) \in \{0, 1\}^{m-n}$ be arbitrarily fixed. For $c \in \{0, 1\}$ let $v(c)$ be the c -successor of v in G_m^C . We have to distinguish two cases.

- v is labelled with some $x_k, 0 \leq k \leq n - 1$. Then, for all $c \in \{0, 1\}$, the c -successor of (v, b) is $(v(c), b(c))$, where $b(c) = (b_0 \oplus L_{k,n} c, \dots, b_{r-1} \oplus L_{k,m-1} c)$.
- v is labelled with some $x_j, n \leq j \leq m - 1$. Then, for all $c \in \{0, 1\}$, if $b_{j-n} \neq c$, the c -successor of (v, b) is the 0-sink. If $b_{j-n} = c$ and $v(c)$ is the *-sink, then let the c -successor of (v, b) be the 1-sink. Otherwise let the c -successor of (v, b) be $(v(c), b)$.

It can be easily checked that R_m (after removing non-reachable nodes) matches all requirements of the Lemma. \square

5 The Main Result

We fix an LFSR-based keystream generator of key-length n with linear bit-stream generator L and a compression function C . We assume that for all

$m \geq 1$ the probability that $C(z)$ is prefix of y for a randomly chosen and uniformly distributed $z \in \{0, 1\}^m$ is the same for all keystreams y . Observe that all generators occurring in this paper have this property. Let us denote this probability by $p_C(m)$.

The cost of our cryptanalysis algorithm depends on two parameters of C . The first is the information rate (per bit) which a keystream y reveals about the first m bits of the underlying internal bitstream. It can be computed as

$$\begin{aligned} \frac{1}{m} I(Z^{(m)}, Y) &= \frac{1}{m} \left(H(Z^{(m)}) - H(Z^{(m)}|Y) \right) = \\ &= \frac{1}{m} (m - \log(p_C(m)2^m)) = -\frac{1}{m} \log(p_C(m)). \end{aligned} \tag{2}$$

where $Z^{(m)}$ denotes a random $z \in \{0, 1\}^m$ and Y a random keystream.

As the compression algorithm computes the keystream in an online manner, the time difference between two succeeding key bits should be small in the average, and not vary too much. This implies the following partition rule: Each internal bit-stream z can be divided into consecutive elementary blocks $z = z^0 z^1 \dots z^{s-1}$, such that $C(z) = y_0 y_1 \dots y_{s-1}$ with $y_j = C(z^j)$ for all $j = 0, \dots, s-1$, and the average length of the elementary blocks is a small constant. This partition rule implies that $p_C(m)$ can be supposed to behave as $p_C(m) = 2^{-\alpha m}$, for a constant $\alpha \in (0, 1]$. Due to (2), α coincides with the information rate of C .

The second parameter of C is the maximal number of output bits which C produces on internal bitstreams of length m . Due to the partition rule, this value can be supposed to behave as γm , for some constant $\gamma \in (0, 1]$. We call γ to be the (best case) compression ratio of C .

Observe that if C always reads the same number k of internal bits for producing one output bit, then $\alpha = \gamma = \frac{1}{k}$. If this number is not a constant then α can be obtained by the formulae

$$2^{-\alpha m} = p_C(m) = \sum_{i=0}^{\lceil \gamma m \rceil} 2^{-i} \text{Prob}_z [|C(z)| = i], \tag{3}$$

where z denotes a random, uniformly distributed element from $\{0, 1\}^m$. Observe that (3) yields $\gamma \geq \alpha$, i.e. $\gamma \alpha^{-1} \geq 1$.

For all $x \in \{0, 1\}^n$ and $m \geq 1$ let $L_{\leq m}(x)$ denote the first m bits of $L(x)$. Note the following design criterion for well-designed keystream generators.

Pseudorandomness Assumption. For all keystreams y and all $m \leq \lceil \alpha^{-1} n \rceil$ it holds that

$$\text{Prob}_z [C(z) \text{ is prefix of } y] \approx \text{Prob}_x [C(L_{\leq m}(x)) \text{ is prefix of } y],$$

where z and x denote uniformly distributed random elements from $\{0, 1\}^m$ and $\{0, 1\}^n$, respectively.

Lemma 2. *If the keystream generator fulfils the above pseudorandomness assumption then for all keystreams y and $m \leq \alpha^{-1}n$ there are approximately $2^{n-\alpha m}$ initial states x for which $C(L_m(x))$ is prefix of y . \square*

Observe that a severe violation of the pseudorandomness assumption implies the possibility of attacking the cipher via a correlation attack. Our main result can now be formulated as

Theorem 2. *Let E be an LFSR-based keystream generator of key-length n with linear bit-stream generator L , and compression function C of information rate α and (best case) compression ratio γ . Let C and L fulfil the BDD- and the pseudorandomness assumption. Then there is an $n^{O(1)}2^{(1-\alpha)/(1+\alpha)n}$ -time bounded algorithm, which computes the secret initial state x from the first $\lceil \gamma\alpha^{-1}n \rceil$ consecutive bits of $y = C(L(x))$.*

As usual, we define the *effective key length* of a cipher of key length n to be the minimal number of polynomial time operations that are necessary to break the cipher. We obtain a bound of $\frac{1-\alpha}{1+\alpha}n$ for the effective key length of keystream generators which fulfil the above conditions.

6 The Algorithm

Let us fix n, E, L, C, α and γ as in Theorem 2. For all $m \geq 1$ let G_m denote the graph ordering defined by C on internal bitstreams of length m . Let y be an arbitrarily fixed keystream which was generated via E . For all $m \geq 1$ let Q_m denote a minimal G_m -FBDD which decides for $z \in \{0,1\}^m$ whether $C(z)$ is prefix of y . Observe that Q_m has to read the first $\lceil \gamma m \rceil$ bits of y . The FBDD-assumption yields that Q_m has polynomial size in m .

For $m \geq n$ let P_m denote the minimal G_m -driven FBDD which decides whether $z \in \{0,1\}^m$ is a linear bitstream generated via L and if $C(z)$ is prefix of y . Observe that by Property 3.03 and Lemma 1

Lemma 3. $|P_m| \leq |Q_m||G_m|^2 2^{m-n}$ for all $m \geq n$. \square

The strategy of our algorithm is simple, it dynamically computes P_m for $m = n, \dots, \lceil \alpha^{-1}n \rceil$. Lemma 2 implies that for $m = \lceil \alpha^{-1}n \rceil$ with high probability only one bit-stream z^* will be accepted by P_m . Due to property 3.04 this bit-stream can be efficiently computed. The first n components of z^* form the initial state that we are searching for.

For all $m \geq n$ let S_m denote a minimal G_m -FBDD which decides for $z = (z_0, \dots, z_m)$ whether $z_m = L_m(z_0, \dots, z_{n-1})$. From Lemma 1 we obtain that $|S_m| \leq 2|G_m|$. Now our algorithm can be formulated as

- (1) $P := Q_n$
- (2) For $m := n + 1$ to $\lceil \alpha^{-1}n \rceil$
- (3) $P := \min(P \wedge Q_m \wedge S_{m-1})$

For the correctness of the minimization in step (3) observe that the definition of G_m implies that G_m is $G_{m'}$ -driven for all $m' \geq m$. It follows from the definitions that for all $m \geq n$ P coincides with P_m after iteration m .

The FBDD-operation $\min(P \wedge Q_m \wedge S_{m-1})$ takes time $p(m)|P_{m-1}|$ for some polynomial p . Consequently, the running time of the algorithm can be estimated by

$$n^{O(1)} \max\{|P_m|, m \geq n\}.$$

Observe that on the one hand, by Lemma 3, $|P_m| \leq p'(m)2^{m-n}$ for some polynomial p' , while on the other hand, by Property 3.02 and Lemma 2, $|P_m| \leq m|One(P_m)|$, where

$$|One(P_m)| \approx 2^{n-\alpha m} = 2^{(1-\alpha)n-\alpha(m-n)}.$$

Consequently, $|P_m|$ does not exceed $n^{O(1)}2^{r(n)}$, where $r(n)$ is the solution of

$$2^{r(n)} = 2^{(1-\alpha)n-\alpha r(n)}$$

which yields $r(n) = \frac{1-\alpha}{1+\alpha}n$. We have proved Theorem 2. \square

7 Applications

We apply Theorem 2 to the keystream generators introduced in section 2. We suppose that these generators fulfill the pseudorandomness assumption, otherwise the running time estimations of our cryptanalysis hold on average. It remains to determine the information rate and the compression ratio, and to prove that the FBDD-assumption is true. For the Connect- k construction it holds $\alpha = \gamma = \frac{1}{k}$. The FBDD-assumption has shown to be true in section 4.

Theorem 3. *For all $k \geq 2$ and all stream ciphers E of key-length n which are a Connect- k construction, our algorithm computes the secret initial state $x \in \{0, 1\}^n$ from the first n bits of $y = E(x)$ in time $n^{O(1)}2^{\frac{k-1}{k+1}n}$. \square*

This is, as far as we know, the best known general upper bound on the effective key-length of the Connect- k construction. Observe that the initial state of one LFSR can be computed efficiently if the initial states of the other LFSRs are known. This leads to a Divide-and-Conquer attack of time $n^{O(1)}2^{\frac{k-1}{k}n}$ which is slightly worse than our result.

For the E_0 -encryption standard in the one-level mode we obtain $\alpha = \gamma = \frac{1}{4}$. The decision if a given internal keystream $z \in \{0, 1\}^m$ yields a prefix of a given keystream y can be computed by π -OBDDs of linear size, see [11], i.e. E_0 fulfils the FBDD-assumption. We obtain

Theorem 4. *For the E_0 -encryption standard with key-length n , our algorithm computes the secret initial state $x \in \{0, 1\}^n$ from the first n bits of $y = E_0(x)$ in time $n^{O(1)}2^{0.6n}$. \square*

Observe that $128 \cdot 0.6 \approx 77$. Note that the best known attacks against the E_0 generator of key length 128 were derived by *Fluhrer* and *Lucks* [7] and *Canniere* [6]. [7] contains a tradeoff result between time and length of available keystream. It varies from 2^{84} necessary encryptions if 132 bit are available to 2^{73} necessary encryptions if 2^{43} bits are available.

Let us now consider the E_0 generator in the two level mode with real key length n and internal key length $N \geq n$. Observe that E_0^2 needs $4 \cdot 4 = 16$ internal bits per key bit for producing the first $N/4$ key bits, while for later key bits only 4 internal bits per key bit are needed. Observe further that our algorithm reaches maximal FBDD-size in iteration $m^* := n + \frac{1-\alpha}{1+\alpha}n$. For $\alpha = 1/16$ this gives $m^* = 32/17n$. As $m^*/16 < N/4$ we obtain $\alpha = \gamma = 1/16$ as relevant parameters for our algorithm on E_0^2 . The decision if a given internal keystream $z \in \{0, 1\}^m$ yields a prefix of a given keystream y can be computed by π -OBDDs of size $O(m)$, where the constant hidden in O is quite large, see [11]. Taking into account that $\frac{1-\alpha}{1+\alpha} = \frac{15}{17} \approx 0.8824$ we get

Theorem 5. *For the E_0^2 -encryption generator with key-length n , our algorithm computes the secret initial state $x \in \{0, 1\}^n$ from the first n bits of $y = E_0^2(x)$ in time $n^{O(1)}2^{0.8824n}$. \square*

As far as we know this is the first nontrivial upper bound on the key length of the E_2^0 generator.

Concerning the self-shrinking generator observe that for all even m and all keystreams y , $shrink(z)$ is prefix of y for exactly $3^{m/2}$ strings z of length m . We obtain an information rate $\alpha = 1 - \log(3)/2 \approx 0.2075$ for the self-shrinking generator by evaluating the relation $2^{-\alpha m}2^m = 3^{m/2}$. The (best case) compression ratio of the self-shrinking generator is obviously 0.5. That the self-shrinking generator fulfils the FBDD-condition was already shown in section 4. Taking into account that for $\alpha = 0.2075$ it holds $\frac{1-\alpha}{1+\alpha} \approx 0.6563$ and $0.5\alpha^{-1} \approx 2.41$ we get

Theorem 6. *For the self-shrinking generator of an n -LFSR L , our algorithm computes the secret initial state $x \in \{0, 1\}^n$ from the first $\lceil 2.41n \rceil$ bits of $y = shrink(L(x))$ in time $n^{O(1)}2^{0.6563n}$. \square*

Observe that the best previously known short-keystream attacks against the self-shrinking generator were given by *Meier* and *Staffelbach* [14] ($2^{0.75n}$ polynomial time operations) and *Zenner et. al.* [19] ($2^{0.694n}$ polynomial time operations). *Mihaljević* [13] presented an attack which yields a tradeoff between time and length of available keystream. It gives $2^{0.5n}$ necessary polynomial time operations if $2^{0.5n}$ bits of keystream are available, and matches our bound of $2^{0.6563n}$ necessary polynomial time operation if $2^{0.3n}$ bits of keystream are available, which is a quite unrealistic assumption.

The difficulty in applying our algorithm to the A5/1 generator is that the compression algorithm reads most of the internal bits twice, one time for the clock control and a certain time later for producing an output key bit. Read-twice BDDs do not have any of the nice algorithmic properties 3.01 - 3.04,

unless $P = NP$. For making the A5/1 generator accessible to our approach we have to modify the keystream generation rule. We define the internal bitstream to be mixed of 6 LFSR-sequences L^0, \dots, L^5 , instead of 3. The first 3 LFSR-sequences are generated by the 3 LFSRs of the A5/1 generator. They are used for producing the output bits. The sequences L^3, L^4, L^5 are used for computing the control values. They are shifted copies of the first 3 sequences, defined by the rules $L_j^{3+k} = L_{j+N_k}^k$, for $k = 0, 1, 2$. As these rules are linear restrictions we get a linear bitstream generated by 6 LFSRs. For this modified version, the decision if a given internal keystream $z \in \{0, 1\}^m$ yields a prefix of a given keystream y can be computed by G_m -FBDDs of size $O(m^3)$. G_m -denotes the graph ordering induced by the clock control. This is the only example where OBDDs do not suffice, we really need FBDDs which allow different variable orderings for different inputs.

The (best case) compression ratio of the modified version of A5/1 is $\gamma = \frac{1}{4}$, as either 4 or 6 new internal bits are used for producing the next output bit. It can be proved that the information rate α is the solution of

$$2^{1-4\alpha} = \frac{1}{4} (3 + 2^{2\alpha}),$$

which yields $\alpha \approx 0.2193$, see [11]. Taking into account that $\frac{1-\alpha}{1+\alpha} \approx 0.6403$ and $\gamma\alpha^{-1} \approx 1.14$ we obtain

Theorem 7. *For an A5/1 generator E of key length n , our algorithm computes the secret initial state x from the first $\lceil 1.14n \rceil$ bits of $y = E(x)$ in time $n^{O(1)}2^{0.6403n}$. \square*

The best previously known short-keystream attack was given by *Golić* [8]. It is against a version of A5/1 generator with keylength 64, which slightly deviates from the specification discovered in [5]. A tight analysis of the time behaviour of *Golić's* attack, when applied to the real A5/1 generator, was given by *Zenner* in [18] and yields 2^{42} polynomial time operations. We get a marginal improvement, as $\lceil 64 \cdot 0.6403 \rceil = 41$. The best long-keystream attacks were given by *Biryukov, Shamir and Wagner* in [2], and *Biham and Dunkelmann* in [1]. After a preprocessing of 2^{42} operations the first attack in [2] breaks the cipher within seconds on a modern PC if around 2^{20} bits of keystream are available. The second attack in [2] breaks the cipher within minutes after a preprocessing of 2^{48} operations and under the condition that around 2^{15} bits of keystream are available. The attack in [1] breaks the cipher within $2^{39.91}$ A5/1 clockings on the basis of $2^{20.8}$ available keystream bits.

8 Discussion

There are classical design criterions for keystream generators like a large period, a large linear complexity, correlation immunity and good local statistics. In this paper we suggest a new one: resistance against FBDD-attacks. We have seen

that there are two strategies to achieve this resistance. The first is to highly compress the internal bitstream (as in the case of E_0^2). This implies a low bit-rate which is not desirable. The second strategy is to design the compression function C in such a way that the decision about the consistence of a given internal bitstream with a given output keystream requires exponential size FBDDs. It is an interesting challenge to look for such constructions. For demonstrating the universality of our approach we presented the FBDD-attack in a very general setting. The obvious disadvantage of this setting is that the algorithm needs a lot of space as all intermediate FBDDs have to be explicitly constructed. It is an interesting open question if the algorithmic idea of FBDD-minimization can be used in a more subtle way for getting, at least for some ciphers, an algorithm which is less space consuming. Another interesting direction of further research is to check whether the FBDD-attack could be successfully combined with other more sophisticated methods of cryptanalysis like the tradeoff attacks suggested in [8], [2] and [1]. Moreover, it would be interesting to clarify by experiments how much do the real sizes of the minimized intermediate FBDDs deviate from the pessimistic upper bounds proved in our analysis.

Acknowledgement

I would like to thank Stefan Lucks, Erik Zenner, Christoph Meinel, Ingo Wegener, Rüdiger Reischuk and some unknown referees for helpful discussions.

References

1. E. Biham, O. Dunkelman. Cryptanalysis of the A5/1 GSM Stream Cipher. Proc. of INDOCRYPT 2000, LNCS 1977, 43-51.
2. A. Biryukov, A. Shamir, D. Wagner. Real Time Cryptanalysis of A5/1 on a PC. Proc. of Fast Software Encryption 2000, LNCS 1978, 1-18.
3. Bluetooth SIG. Bluetooth Specification Version 1.0 B, <http://www.bluetooth.com/>
4. R. E. Bryant. Graph-based algorithms for Boolean function manipulations. IEEE Trans. on Computers 35, 1986, 677-691.
5. M. Briceno, I. Goldberg, D. Wagner. A pedagogical implementation of A5/1. <http://www.scard.org>, May 1999.
6. C. de Canniere. Analysis of the Bluetooth Stream Cipher. Master's Project COSIC, Leuven, 2001.
7. S. R. Fluhrer, S. Lucks. Analysis of the E_0 Encryption System. Technical Report, Universität Mannheim 2001.
8. J. D. Golić. Cryptanalysis of alleged A5/1 stream cipher. Proc. of EUROCRYPT'97, LNCS 1233, 239-255.
9. J. Gergov, Ch. Meinel. Efficient Boolean function manipulation with OBDDs can be generalized to FBDDs. IEEE Trans. on Computers 43, 1994, 1197-1209.
10. S. W. Golomb. Shift Register Sequences. Aegean Park Press, Laguna Hills, revised edition 1982.
11. M. Krause. BDD-based Cryptanalysis of Keystream Generators. Report 2001/092 in the Cryptology ePrint Archive (<http://eprint.iacr.org/curr/>).

12. Ch. Meinel. Modified Branching Programs and their Computational Power. LNCS 370, 1989.
13. M. J. Mihaljević. A faster Cryptanalysis of the Self-Shrinking Generator. Proc. of ACIPS'96, LNCS 1172, 182-189.
14. W. Meier, O. Staffelbach. The Self-Shrinking Generator. Proc. of EUROCRYPT'94, LNCS 950, 205-214.
15. R. A. Rueppel. Stream Ciphers. Contemporary Cryptology: The Science of Information Integrity. G. Simmons ed., IEEE Press New York, 1991.
16. D. Sieling, I. Wegener. Graph driven BDDs - a new data structure for Boolean functions. Theoretical Computer Science 141, 1995, 283-310.
17. I. Wegener. Branching Programs and Binary Decision Diagrams. SIAM Monographs on Discrete Mathematics and Applications. Philadelphia 2000.
18. E. Zenner. Kryptographische Protokolle im GSM Standard: Beschreibung und Kryptanalyse (in german). Master Thesis, University of Mannheim, 1999.
19. E. Zenner, M. Krause, S. Lucks. Improved Cryptanalysis of the Self-Shrinking Generator. Proc. of ACIPS'2001, LNCS 2119, 21-35.